

MedConsensus: A Secure AI-Assisted Multi-Doctor Validation Framework

Specially for Cancer Patients

Vaishnavi Pandurang Belote
Department of Computer Science
Dr. D. Y. Patil Arts, Commerce & Science College
Pimpri, Pune, Maharashtra, India

Abstract - Telemedicine platforms have improved access to healthcare, yet delayed physician responses, clinical workload imbalance, and dependency on single doctor decisions continue to pose risks to patient safety. This paper proposes an AI assisted Multi Agent Multi Doctor Consensus Framework for secure and real time telemedical decision support. The proposed system integrates a human in the loop artificial intelligence triage engine with a weighted multi doctor validation mechanism to reduce single point of failure in clinical decision making. An anonymization layer is incorporated to protect patient identity through encrypted storage and masked identifiers, while licensed doctor verification ensures platform authenticity and trust.

The architecture follows a privacy by design approach using a microservices based infrastructure. The implementation is technically feasible using Python for natural language processing, TensorFlow for predictive modeling, React for frontend interface, Node or Django for backend services, and secure databases with encryption standards. Role based access control, token based authentication, and secure cloud deployment support scalability and reliability.

The framework introduces weighted clinical consensus where specialist experience and credential relevance influence validation strength, while artificial intelligence remains strictly decision support and not a replacement for physicians. Simulated evaluation indicates potential reduction in response time, improved reliability, and optimized workload distribution. The proposed model contributes toward safer, faster, and more trustworthy digital healthcare systems.

Keywords - Artificial Intelligence in Healthcare, Telemedicine Systems, Multi Doctor Consensus Framework, Clinical Decision Support System, Human in the Loop AI, Privacy by Design Architecture, Health Data Security, Physician Workload Optimization, Anonymized Medical Communication, Weighted Validation Model, Secure Telehealth Infrastructure

I. INTRODUCTION

A. Background

The rapid growth of telemedicine has transformed healthcare delivery by enabling remote consultation, digital prescriptions, and virtual monitoring. Global adoption of telehealth systems increased significantly due to technological advancements, internet accessibility, and the demand for convenient medical services. At the same time, artificial intelligence has emerged as a powerful tool in healthcare, supporting clinical decision making, disease prediction, symptom classification, and workflow optimization.

Despite these advancements, healthcare systems continue to experience rising physician workload and burnout. Doctors often manage high patient volumes, administrative responsibilities, and time pressure, which can delay response times in teleconsultation platforms. In critical scenarios, delayed medical guidance may increase the risk of complications. Therefore, there is a growing need for intelligent systems that support faster, safer, and more reliable medical decision making while maintaining physician authority.

B. Problem Statement

Although telemedicine improves accessibility, several limitations remain unresolved. Delayed responses from physicians can negatively impact patient outcomes, especially in time sensitive cases. Excessive workload and uneven case distribution contribute to inefficiency and stress among healthcare professionals.

Additionally, conventional systems typically depend on a single doctor's opinion, creating a single point of failure in clinical decision making. The absence of multi doctor validation increases variability in diagnosis and treatment recommendations. Privacy concerns also persist, as patient identity information may be exposed or insufficiently protected. Furthermore, verification of medical professionals remains a challenge in some digital platforms, increasing the risk of unauthorized or unqualified practitioners.

These limitations highlight the necessity for a structured, secure, and collaborative telemedical framework.

C. Research Gap

Existing telemedicine platforms primarily focus on one to one consultation models and basic artificial intelligence assistance. However, there is limited research on structured integration of artificial intelligence with a weighted multi doctor consensus mechanism for clinical validation. Current systems do not commonly incorporate formalized doctor weight allocation based on experience and specialization.

Moreover, many platforms lack a clearly defined anonymization layer that separates patient identity from clinical data using a privacy by design approach. Real time global workload balancing across time zones and availability levels is also insufficiently addressed in current telehealth infrastructures.

These gaps indicate the need for a comprehensive framework that combines artificial intelligence, structured

multi doctor validation, identity protection, and dynamic workload optimization.

D. Objective

The primary objective of this research is to design a secure artificial intelligence assisted multi doctor consensus platform for real time telemedical decision support. The proposed framework aims to reduce response time, enhance reliability, and minimize dependency on single doctor evaluation. It also seeks to protect patient identity through anonymization and encrypted storage mechanisms while ensuring strict verification of licensed medical professionals.

The system is designed to support physicians rather than replace them, improving decision efficiency while preserving clinical authority.

E. Hypothesis

This research hypothesizes that a weighted multi doctor artificial intelligence assisted framework can significantly improve patient safety, reduce response delays, and decrease decision variability compared to conventional single doctor telemedicine systems. By integrating structured consensus validation, privacy by design architecture, and workload distribution mechanisms, the proposed model is expected to enhance trust, reliability, and operational efficiency in digital healthcare environments.

II. METHODOLOGY

A. System Architecture

The proposed framework follows a layered and modular architecture designed using a microservices approach to ensure scalability, security, and maintainability. The system consists of the following major layers:

Patient Interface Layer:

This layer allows patients to submit symptoms, upload prescriptions or reports, and receive validated medical guidance. The interface is designed to be simple, secure, and accessible across web and mobile platforms.

Anonymization Layer:

Before any clinical data is processed, personally identifiable information is separated from medical content. Patient names, contact details, and identifiers are replaced with system-generated anonymized IDs. This layer ensures privacy by design and prevents direct exposure of identity to consulting doctors.

AI Triage Engine:

The anonymized clinical data is processed through a natural language processing based triage engine that categorizes symptoms, estimates severity level, and generates preliminary suggestions with a confidence score.

Doctor Routing Engine:

Cases are dynamically assigned to available doctors based on specialization, experience, time zone, and workload distribution parameters.

Multi Doctor Validation Layer:

At least two to three licensed doctors review the AI assisted case summary and provide independent opinions.

Weighted Consensus Engine:

Doctor validations are processed through a structured weighted scoring mechanism to generate a final consensus recommendation.

Secure Identity Vault:

Encrypted storage maintains patient personal data separately from clinical data using secure key management protocols.

Audit and Logging Module:

All actions are logged to ensure traceability, transparency, and accountability.

The architecture ensures separation of concerns and reduces dependency on any single component.

B. Technical Requirements

The system is implemented using a modern, scalable technology stack.

Programming Stack:

Python is used for artificial intelligence development and natural language processing. TensorFlow or PyTorch supports predictive modeling. The frontend interface is developed using React.js for responsive interaction. Backend services are implemented using Node.js or Django frameworks with RESTful APIs for communication. PostgreSQL or MongoDB is used for secure data storage.

Security Stack:

Data encryption is implemented using AES based encryption standards. Secure communication is ensured through HTTPS with SSL or TLS protocols. JWT based authentication manages session validation. Role Based Access Control restricts data access according to user roles. Two factor authentication enhances login security.

Infrastructure:

Cloud deployment is supported on platforms such as AWS, Azure, or GCP. Docker containers are used for environment consistency, while Kubernetes enables orchestration and automatic scaling. Load balancers distribute incoming requests efficiently.

The microservices architecture allows independent scaling of AI, routing, and validation components.

C. Identity Protection Model

The proposed system follows a privacy by design principle. Personally identifiable information is stored separately in an encrypted identity vault. Doctors only access anonymized patient IDs and relevant clinical details required for evaluation. A zero knowledge access model ensures that no single user can access both identity and clinical data simultaneously without authorization. Every interaction is recorded through an audit trail system. Doctor credentials undergo periodic verification to maintain platform integrity.

The framework is conceptually aligned with global privacy standards such as HIPAA style protection and GDPR aligned data minimization. Patient consent is required before any data processing or consultation begins.

D. Weighted Multi Doctor Consensus Mechanism

The core innovation of the proposed framework is the weighted multi doctor consensus model. Each case requires

validation from a minimum predefined number of licensed doctors.

Doctor contributions are assigned weights based on parameters such as years of experience, specialty relevance, and credential level. A specialist in the relevant domain carries higher influence compared to a general practitioner. Senior physicians contribute higher validation weight than junior practitioners. The artificial intelligence output carries minimal influence and functions strictly as decision support.

A structured scoring algorithm aggregates individual doctor confidence scores to compute a final consensus result. Conflict detection mechanisms identify contradictory recommendations and may trigger additional review.

This structured validation approach reduces decision variability and minimizes the risk of single point failure in medical consultation.

E. AI Triage Model

The AI triage engine applies natural language processing techniques to analyze patient reported symptoms and uploaded medical documents. The model performs symptom classification, severity detection, and risk probability estimation.

A recommendation summary is generated along with an AI confidence score. However, the system strictly enforces a human in the loop model, where artificial intelligence does not produce final diagnosis. Final medical recommendations are always determined by licensed physicians.

This design ensures responsible AI deployment and preserves clinical authority.

F. Global Workload Balancing

To reduce physician overload, the system incorporates a dynamic workload balancing mechanism. Case assignment considers doctor availability status, time zone alignment, specialization match, and current queue length.

If a doctor becomes inactive or fails to respond within a defined interval, automatic reassignment is triggered. This dynamic load distribution improves response time and contributes to physician burnout reduction.

The routing engine continuously monitors system performance to maintain balanced case distribution across regions.

G. Figures/Diagrams

Fig 1: Platform Architecture

[USER ACCESS LAYER]

Patients
Secure Login

Doctors
License Based Login

Admin
System Control

↓

[ANONYMIZATION LAYER]

Extract Personal Identifiable Information
 Encrypt Sensitive Data
 Store in Secure Identity Vault
 Generate Anonymous Patient ID
 Forward Only Clinical Data

↓

[AI TRIAGE ENGINE]

Symptom Text Processing
 NLP Based Classification
 Severity Detection
 Risk Probability Score
 Generate AI Suggestion Decision Support Only

↓

[DOCTOR ROUTING ENGINE]

Match Required Specialty
 Check Doctor Availability
 Time Zone Based Allocation
 Workload Optimization
 Assign Case to Two or Three Doctors

↓

[MULTI DOCTOR VALIDATION LAYER]

Independent Clinical Review
 Approve Modify or Provide New Suggestion
 Submit Confidence Score

↓

[WEIGHTED CONSENSUS ENGINE]

Calculate Doctor Weights
 Experience into Specialty into Credential Level
 Aggregate Validation Scores
 Detect Disagreements
 Escalate to Senior Specialist if Required
 Generate Final Verified Decision

↓

[FINAL RESPONSE OUTPUT]

Consensus Based Medical Recommendation
 Doctor Name Degree and Hospital Visible
 Audit Log Stored for Compliance

Fig 2: Interfaces

[PATIENT INTERFACE]

Secure Registration
 Consent Acceptance
 Symptom Entry Panel
 Upload Medical Reports
 View AI Severity Level
 Track Case Status
 Receive Final Consensus Recommendation
 View Approved Doctor Name Degree Hospital
 Follow Up Instructions



[DOCTOR INTERFACE]

Verified Login with License Authentication
 View Anonymous Patient ID Only
 Review Symptoms and Reports
 View AI Suggestion and Risk Score
 Approve Modify or Reject Suggestion
 Submit Clinical Decision
 Provide Confidence Score
 View Assigned Cases and Workload



[ADMIN INTERFACE]

Verify Doctor Credentials
 Approve or Reject Registration
 Monitor License Revalidation
 View Audit Logs
 Monitor System Security
 Oversee Workload Distribution
 Manage Conflict Escalation
 Monitor Encryption and Compliance

III. RESULTS

A controlled prototype evaluation of the proposed AI-Assisted Multi-Doctor Consensus Framework was conducted using simulated outpatient case scenarios. The analysis focused on response time behavior, AI triage stability, consensus agreement, conflict handling, workload distribution, and system throughput.

A. Response Time Observation

The AI triage engine initiated processing immediately after patient submission. Case routing to qualified doctors occurred automatically. Compared to traditional manual review systems, the proposed framework demonstrated faster initiation of medical evaluation and reduced waiting intervals before receiving validated responses.

Early AI-based structuring of cases minimized initial processing delay.

B. AI Classification Stability

The NLP-based triage model consistently categorized symptom severity and generated structured summaries for doctor review. The AI acted strictly as a decision-support system, while final diagnosis remained under licensed physicians.

Human-in-the-loop control preserved clinical authority and safety.

C. Multi-Doctor Consensus Behavior

In most simulated cases, doctors reached aligned clinical recommendations. In situations of disagreement, the weighted consensus mechanism aggregated inputs based on experience and specialty relevance. Complex cases were escalated to senior specialists.

Reduction of single-point-of-failure through structured multi-doctor validation.

D. Workload Distribution and Throughput

The routing engine allocated cases according to specialty match and active load. Automatic reassignment prevented backlog when a doctor was inactive. The microservices architecture handled multiple concurrent submissions without instability.

Dynamic load distribution supported physician workload optimization.

Parameter	Traditional Model	Proposed Framework
Initial Processing	Manual review	Immediate AI triage
Decision Authority	Single doctor	Multi-doctor consensus
Risk Handling	Informal backup	Weighted escalation model
Identity Exposure	Visible patient data	Anonymized ID system
Workload Management	Manual allocation	Automated load balancing
AI Role	Limited chatbot	Structured decision support
Accountability	Basic logging	Integrated audit tracking

IV. DISCUSSION

The proposed AI-Assisted Multi-Doctor Consensus Framework introduces a structural shift from traditional telemedicine models that typically depend on a single physician's evaluation. In conventional systems, decision-making relies on individual clinical judgment, which may be influenced by time pressure, workload, or limited case context. The proposed framework addresses this limitation by integrating a weighted multi-doctor validation mechanism supported by AI-based triage.

A. Comparison with Traditional Telemedicine

Traditional telemedicine platforms generally follow a single-doctor consultation approach, where one physician reviews patient data and provides recommendations. While efficient for routine cases, this model presents a potential single-point-of-failure risk. In contrast, the proposed framework introduces structured multi-doctor consensus, automated routing, and privacy-preserving identity masking. This layered approach enhances reliability and strengthens clinical governance.

A key contribution of this model is the reduction of clinical decision variability, achieved through aggregated expert validation rather than isolated opinion.

B. Benefits of Consensus-Based Validation

The weighted consensus mechanism improves diagnostic confidence by incorporating multiple expert perspectives. Senior specialists contribute proportionally higher influence in complex cases, ensuring medical hierarchy is respected. Conflict detection and escalation modules further enhance safety in uncertain scenarios. This structure reduces dependency on a single physician and promotes balanced decision-making.

C. AI as Decision Support, Not Replacement

The AI triage engine functions strictly as a decision-support component. It structures patient inputs, categorizes severity, and assists in routing, but does not issue final diagnoses. All medical decisions remain under licensed physician control, maintaining ethical compliance and professional accountability.

D. Ethical and Privacy Considerations

The integration of an anonymization layer and encrypted identity vault supports a privacy-by-design architecture. Patient-identifiable information is protected through secure storage and controlled access mechanisms. The system promotes transparency through audit logging and traceability while maintaining consent-based data handling principles.

E. Scalability and Cost Considerations

The microservices-based cloud infrastructure allows horizontal scaling according to user demand. Automated routing reduces administrative overhead, while distributed

architecture supports global deployment. Although initial implementation may require investment in secure infrastructure and verification systems, long-term operational efficiency may reduce repetitive workload and improve consultation throughput.

F. Limitations

Despite its advantages, the framework has certain limitations. The effectiveness of consensus validation depends on doctor availability across time zones. Regulatory approval and jurisdiction-specific compliance may affect deployment. Additionally, AI-driven triage systems require continuous monitoring to prevent potential bias or misclassification.

V. CONCLUSION

This research addresses critical limitations in conventional telemedicine systems, including delayed response times, physician overload, privacy vulnerabilities, and the risks associated with single-doctor dependency. By identifying the single-point-of-failure problem in traditional consultation models, this study proposes a structured and secure alternative framework.

The core contribution of this work is the development of an AI-assisted multi-agent, multi-doctor weighted consensus mechanism that enhances clinical reliability while preserving human authority in decision-making. The integration of a privacy-first anonymization layer, encrypted identity vault, and role-based access control ensures secure handling of sensitive medical data. Unlike automated diagnostic systems, the proposed architecture maintains a human-in-the-loop model, where AI supports but does not replace licensed medical professionals.

Through structured validation, intelligent routing, and workload balancing, the framework demonstrates improved consistency in clinical recommendations and optimized response workflows. The reduction of isolated decision dependency strengthens system trust and promotes safer telemedical practice.

Overall, the proposed framework presents a scalable, privacy-aware, and ethically responsible model for next-generation telemedicine platforms. With further validation and regulatory alignment, this approach has the potential to significantly influence the future design of secure, collaborative, and intelligent digital healthcare systems.

VI. FUTURE WORK

The current research establishes a conceptual and architectural foundation for a multi-doctor consensus platform. Future developments will focus on the following key areas:

Blockchain Integration for Record Validation: Implementing a decentralized ledger system to create immutable records of clinical validations. This will ensure that medical history and doctor approvals are tamper-proof and fully auditable.

Electronic Health Record (EHR) Integration: Developing standardized APIs to synchronize the platform with existing hospital EHR systems. This will allow doctors to view a patient's historical medical data instantly for more accurate consensus building.

Emergency Prioritization Module: Enhancing the AI triage engine to detect "Red Flag" keywords (e.g., chest pain, difficulty breathing) to automatically bypass the standard routing queue and trigger immediate emergency alerts for the patient.

Self-Learning AI Optimization: Utilizing federated learning techniques to improve the AI Triage model's accuracy over time. This allows the system to learn from doctor corrections without compromising patient privacy or data security.

Government and Public Health Integration: Collaborating with national health departments to integrate the framework into public hospital networks, specifically to support rural areas with limited specialist access.

Cross-Platform Mobile Deployment: Developing native iOS and Android applications featuring secure biometric authentication (FaceID/Fingerprint) to make the platform accessible for real-time patient monitoring and doctor consultation.

VII. ACKNOWLEDGEMENT

I would like to express my sincere gratitude to Dr. D Y Patil ACS College and Organizers, for providing me with the invaluable opportunity to undertake this research project. Their guidance and support were crucial in the completion of this work.

VIII. REFERENCES

- [1] WHO, "Global Strategy on Digital Health 2020-2025," WHO Publication, 2021.
- [2] B. Shanafelt, "Physician Burnout: Causes & Cures," Mayo Clinic Proc., 92(1), 2017.
- [3] Cavoukian, "Privacy by Design," IPC Ontario, 2009.
- [4] HHS, "HIPAA Privacy Rule," HHS.gov, 2003.
- [5] EU, "GDPR Regulation (EU) 2016/679," Official Journal of EU, 2016.
- [6] ICMR, "Ethical Guidelines for AI in Healthcare," ICMR.gov.in, 2023.
- [7] S. Reddy, "AI in Healthcare: Review of Triage Systems," IEEE Access, vol. 8, 2020.
- [8] L. Wang, "Multi-Agent Systems for Healthcare," Expert Syst. Appl., vol. 145, 2020.
- [9] Sullivan, "Weighted Consensus Algorithms for Medical Validation," Int. J. Med. Inform., 112, 2018.
- [10] NAM, "AI in Health Care: Promise and Peril," NAM Special Pub, 2019.
- [11] Meyer, "Zero-Trust Architecture for Medical Identity," IEEE J. Biomed. Health Inform., 25(3), 2021.
- [12] M. West, "Dynamic Load Balancing in Telemedical Routing," Softw. Pract. Exp., 51(8), 2021.